A Physics/Engineering of Failure Based Analysis and Tool for Quantifying Residual Risks in Spaceflight Hardware

Steven L. Cornford, Mark Gibbel, Martin Feather, David Oberhettinger
• Jet Propulsion Laboratory, California Institute of Technology
• Pasadena, California

SUMMARY & CONCLUSIONS

NASA is supporting efforts to improve the verification and validation process and the risk management process for spaceflight projects. A physics-of-failure based Defect Detection and Prevention (DDP) methodology has been developed and is currently being implemented on various NASA projects and as part of NASA's new model-based spacecraft development environment. DDP weights the criticality of the various relevant FM's by including the likelihood and impact on mission requirements.

The methodology begins with prioritizing the risks (or FM's/mechanisms) (FM's) relevant to a mission which need to be addressed. These risks can be detected or prevented through the implementation of a set of mission assurance activities—referred to herein as "PACTs.[1] Each of these PACTs has some effectiveness against one or more FM's but also has an associated resource cost. The FM's can be weighted according to their likelihood of occurrence and their mission impact should they occur. The net effectiveness of various combinations of PACTs can then be evaluated against these weighted FM's to obtain the residual risk for each of these FM's and the associated resource costs to achieve these risk levels. The process thus identifies the project-relevant "tall pole" FM's and design drivers and allows real time tailoring with the evolution of the design and technology content. The DDP methodology allows risk management in its truest sense: it identifies and assesses risk, provides options and tools for risk decision making and mitigation and allows for real-time tracking of current risk status.

1. INTRODUCTION

NASA continues to make progress in response to its mandate to fabricate and operate spacecraft "faster, better, and cheaper".[2] The posture of risk avoidance has given way to active risk management. A key element of NASA's risk management approach is to consider "risk as a resource".[3] Like schedule, mass and power, risk is now a resource to

---

[1] PACTs: **P**reventions (typically design measures), **A**nalyses, process **C**ontrols (e.g., parts selection), and **T**ests

[2] *NASA Strategic Management Handbook,* National Aeronautics and Space Administration, Washington D.C., April 1996.

[3] M. A. Greenfield and T. E. Gindorf, *Risk as a Resource – A New Paradigm,* Proceedings of the ESREL 96 - PSAM III Conference, Create, Greece, Vol. 3, pp. 1597, Springer-Verlag, Berlin, June 24 - 28, 1996.

be traded against other resources and optimized subject to constraints. This process has been facilitated by the NASA focus on developing better risk management tools and methods.[4]

The typical NASA project is evolving to an 18 month development cycle in which decisions will need to be made near "real-time" in a model-based development environment. Like any decision, there is a chance that it will be wrong. This leads to the chance that something will fail. This is where risk is introduced, "What is the risk of failure due to a given decision or action?" Since risk can be reduced by expending resources, and risk itself is a resource, an integrated methodology for trading these resources would be valuable.

Early decisions usually have the most influence on the project risk, but the realities of the fast-track spacecraft development cycle often necessitate decisions based on incomplete data. Furthermore, each of these decisions may result in different, or additional, derived requirements. The DDP process allows the requirements to be captured, the current risks to be estimated, and tradeoffs to be made with the available data. These data regarding risks and consequences ranges from engineering judgement to actual flight article test data depending on the stage of the project development process. DDP integrates this variety of data in a "top down" approach which is synchronized to the project development cycle.

1.1 *Definitions*

The following technical terms are defined as they are used in the context of this paper:

Escape: A situation where a FM escapes detection by a PACT and consequently must be addressed by a subsequent PACT. Escapes may result from an inadequate application of the PACT (e.g., test duration too short) or because the PACT is incapable of detecting the FM (e.g., detecting excessive mirror surface roughness by means of a vibration test).

Failure: An incident in which a circuit or subsystem does not perform an intended function. This may range from reduced gain of 1dB, to an explosion.

FM: The characteristic manner in which a failure occurs, independent of the reason for failure; the condition or state which is the end result of a particular failure mechanism. In the DDP process, the FM's are arranged in a tree structure.

Mission/Project Requirements: A set of characteristics or distinguishing features that are needed to meet operational needs and comply with applicable policy and practices. In the DDP process, the requirements are grouped in a tree structure.

---

[4] Liam Sarsfield, *"The Cosmos on a Shoe String"*, RAND Critical Technologies Institute, Washington D.C., MR-864-OSTP, 1998.

PACTs: An acronym for "preventative measures, analyses, process controls and tests," PACTs are the collection of possible prevention and detection activities. As a product element passes through a PACT (e.g. a test), anomalies (FM's) are observed (detected) and presumably fixed.

Tall Poles: FM's or failure mechanisms that stand out from the others because of either the likelihood of their occurrence or the mission impact should they occur.

2. *METHODOLOGY*

2.1 Overview

The application of the DDP process involves four steps: develop the requirements matrix, develop the effectiveness matrix, optimize the residual risk (subject to constraints), and iterate throughout the project life cycle. These four steps, shown in Figure 1, involve creating and populating two matrices, performing a tailoring step, and iterating the process as requirements and risk evolve. The key inputs to the DDP process are the FM's (or failure mechanisms at lower levels), mission/project requirements (which may be

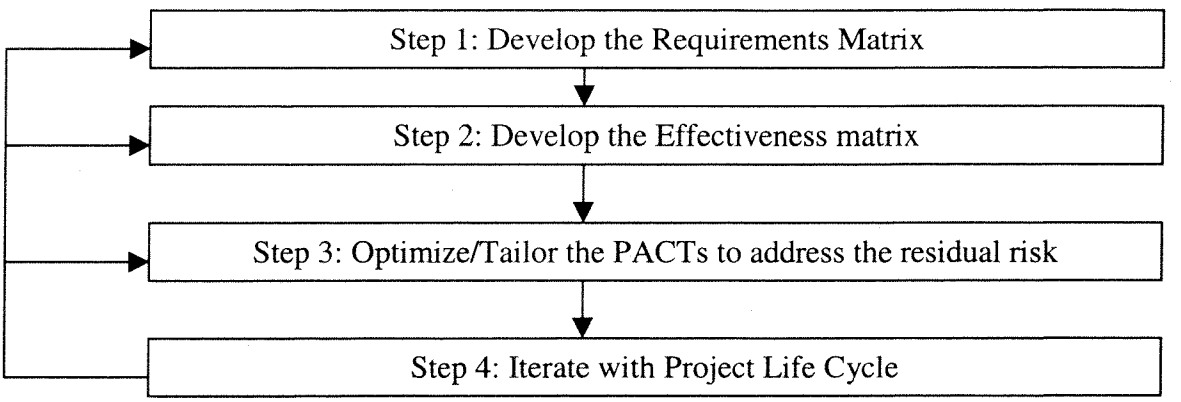

Figure 1. DDP is a Four-Step Process

derived sub-system requirements at lower levels), and the suite of available PACTs.

In Figure 2, each box represents a collection of mission assurance measures, and the dotted lines represent "escapes"-- FM's that were not detected or prevented by the engineering activity. The optimal mission assurance program is one that identifies and retains the minimum set of measures necessary to expose a critical failure or prevent it from occurring in flight.

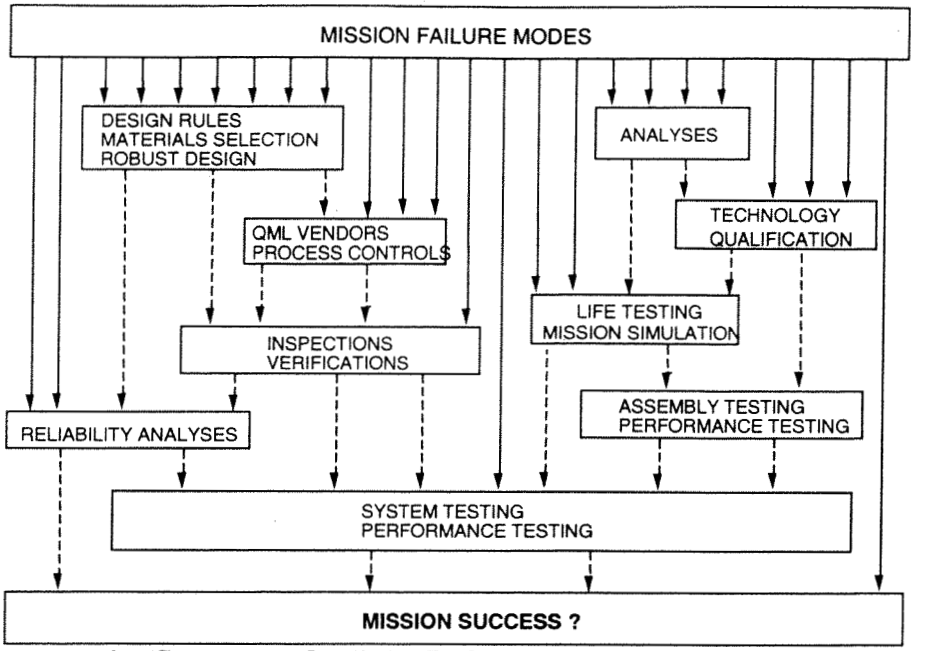

Figure 2. "Screening Out" the Defects (illustrative diagram—not to scale)

2.2 Four-Step Process

The DDP methodology consists of four steps. An ongoing example will be used to clarify the process. For simplicity this discussion will be constrained to one level of hierarchy.

Step 1: Develop the Requirements Matrix. This step actually involves 3 sub-steps:

1. Identify the requirements (rows of the matrix),
2. Identify the FM's (columns of the matrix), and
3. Populate the matrix (the matrix elements).

This produces a prioritized set of FM's (risk elements) in which the "tallest pole" is (loosely) the FM which has the greatest impact on the most important requirements, and the "shortest pole" is the FM which has the least impact on the least important requirements. This prioritized set allows mission assurance (and project) personnel to focus their attention on a prioritized list of risk elements. Let us now examine in more detail how one achieves this critical result.

Step 1.1: Identify Requirements. This first step requires involvement from project personnel (the customers) and entails the identification, weighting and grouping of the requirements for the program or project under evaluation. This grouping may be by requirement type, or by the various instrument requirements, etc. The requirements are also weighted by the relative importance to the project. As an obvious example, the secondary mission requirements are weighted less than the primary mission requirements.

Tool steps: To input requirements, select the "Rqmt" window, use the pull-down menus to select "Node -> Add Node" and then enter in the requirement and it's associated attributes. Generate trees as logical.

Example: A simple example illustrates a grouping into primary and secondary mission requirements. Under primary mission requirements, the weightings are all 10 except for the launch date, which is weighted as 8. Under secondary mission requirements, the weightings are 2 and 3, respectively.

Sample DDP output

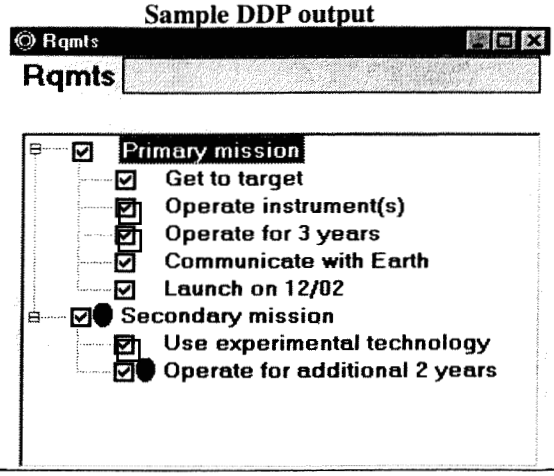

Step 1.2: Identify FM's. Next, the potentially relevant FM's (or risk elements) are identified and grouped. FM's may be identified by a variety of techniques including brainstorming, interviews with design engineers, reviews of lessons learned, incorporating existing fault trees or FM's, effects, and criticality analyses (FMECAs), etc.

This identification and grouping is best done at a level consistent with the requirements. For example, the requirement "Operate for 3 Years" leads one to worry about "Life Issues". As the details of the design emerge, and the DDP process continues to stay synchronized with the project life cycle, the FM's start to become more specific and "Life Issues" is then broken into fatigue, expended consumables, environmental degradation, etc. This process is most easily and efficiently accomplished using a "critical mass" of experts since each discipline expert usually has a unique perspective, and experience indicates that many of the risk elements lie in the interfaces into which one person rarely has complete insight.

Tool steps: To input FM's, select the "FMs" window, use the pull-down menus to select "Node -> Add Node" and then enter in the FM and it's associated attributes. Generate trees as logical.

Example: A simple example illustrates a collection of FM's and some of the subordinate FM's. These FM's represent the collection of possible risk elements and will later be prioritized.

Sample DDP output

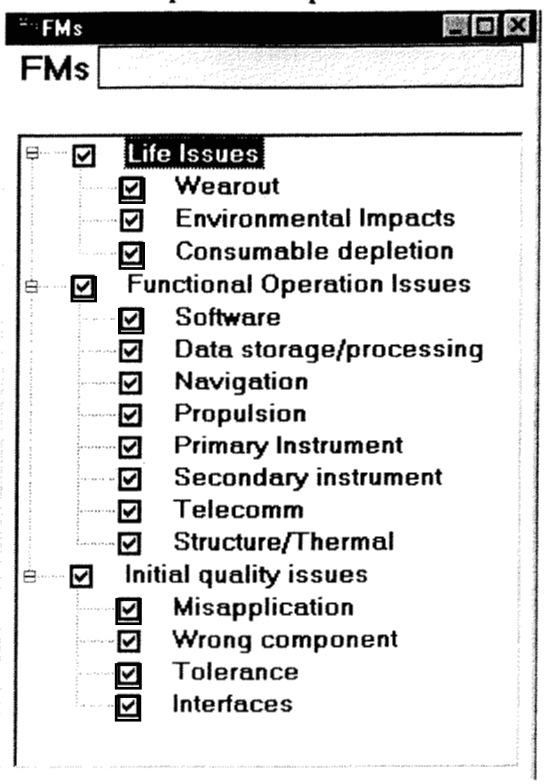

Step 1.3: Populate the Requirements Matrix. Now that the requirements and FM's have been identified, one can begin to evaluate the impact of each FM on each requirement. The default approach is to evaluate the percentage of each requirement lost should the FM occur. At the higher levels of evaluation (such as those in this example), it is recommended to use a Taguchi-type non-linear scale such as 0, 1,3, and 9. This is accomplished by entering 0, 0.1, 0.3 and 0.9 representing the fraction of the requirement impacted by the FM.

Example: Using the previously identified FM's and requirements, the Requirements matrix is generated (partially shown). The totals under the FM's represent the total impact and are used to perform the 'tall pole' assessment. The totals next to the requirements can be used to identify "driving requirements", such as Operate for 3 Years.

Sample DDP output

RxFM

RxFM 0 or empty = none lost; 1 = 100% lost

		FMs				Functional Operation Is		
		FMs	Wearout	Environm Impacts		Software	Data storage,	Ne
Rqmts	Rqmts	Totals	29.467	56.533		81.933	65.867	4
Primary mission	Get to target	118.33	0.1	0.7		1	0.3	
	Operate instrument(s)	121.67	0.1	0.7		1	1	
	Operate for 3 years	130	0.7	0.9		0.3	1	
	Communicate with Earth	115	0.1	0.3		1	0.3	
	Launch on 12/02	100				0.7	1	
	Use experimental technology	48.8	0.1	0.3		0.7	0.1	

Tool steps: To input impacts, select the "RxFM" window, and left-click on the intersection of the requirement and FM for which data is to be entered. Complete the pop-up window entries, and go on the next intersection.

Outputs of Step 1: One output of this step is a list of driving requirements. That is, the horizontal sum (for each requirement across all FM's) of matrix entries identifies the total extent to which a given requirement is "'at risk" from the FM's. Since each requirement (and it's weight) ultimately drives the risk element priority, this output gives the project a chance to see which requirements were too aggressive or add no value. Note that generally any relaxation of requirements changes the FM impacts, which in turn reduces, or redistributes, the initial risk balance ("tall poles").

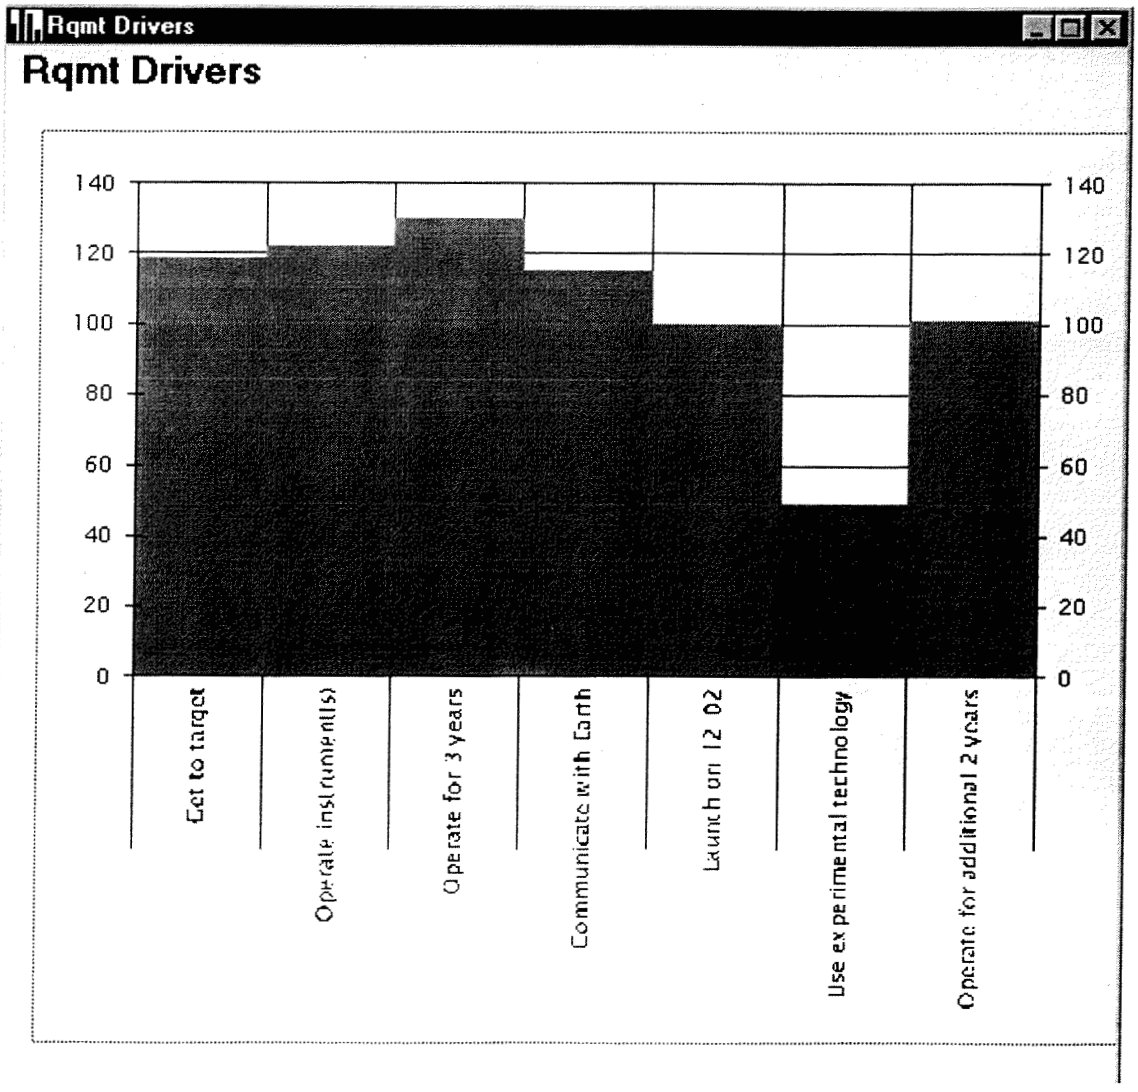

Figure 3. "Requirement drivers" for the example in the text. Note that the requirements most likely to be impacted by the risk elements can be easily identified.

As discussed above, Step 1 also results in a prioritized list (or Pareto diagram) of risk elements which can be used by various project personnel to focus their work. This collection of weighted risk elements is also the key input to the next step. These represent the project risks and must be reduced to the desired levels.

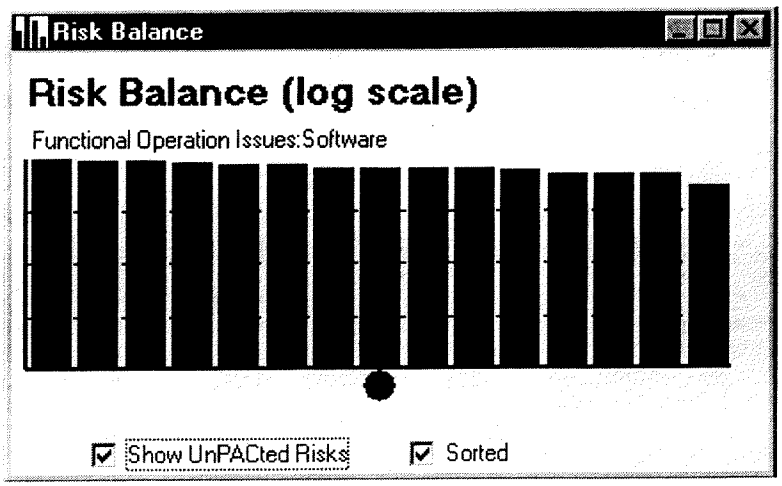

Figure 4 Initial risk balance for the example in the text. Note that no PACTs have been selected and the FMs are ordered according to total impact. The cursor (not visible) location is on the "tallest pole," which happens to be "Functional Operation Issues: Software" as the pop-up text indicates.

Step 2: Develop the Effectiveness Matrix. Now that we have identified and prioritized the relevant risk issues, we can begin to explore possibilities for preventing or detecting them. This step really involves only two sub-steps since the FM's (columns) have already been identified in the previous step. These two sub-steps are (1) identifying detection and prevention options (PACTs) and (2) evaluating their effectiveness against the identified FM's. Obviously the more PACTs we do, the more we lower the risk, but each PACT has an associated resource cost (e.g., primarily mass for radiation shielding, but cost and schedule for radiation testing). There are thus optimal combinations of PACTs that fit the project resource constraints. Completion of the Effectiveness Matrix puts us in the position to tailor these activities.

Step 2.1: Identify the PACT Options. The list of PACT options is both long, and "pre-canned," in the DDP tool. Many of the usual PACTs at assembly level and above are very similar from project to project although they may be heavily tailored. The usual PACTs are already included in the tool (e.g., assembly-level thermal vacuum testing, system-level random vibration testing). Other PACTs are can be expressed generically, but can be made more project specific. For example, a project without optics would delete "Optical Testing" as a candidate PACT, while a project with optics would replace "Optical Testing" with something such as "Optical Alignment Testing" and "Detector Calibration Testing."

The DDP tool is intended to assist the user in getting all of the PACT options "on the table" with names recognized by project personnel. In the early stages of the DDP process, it is better to have too many PACTs than too few, since the following steps will select or un-select the PACTs to see what subset produces the optimal solution.

Tool steps: To input PACTs, select the "PACTs" window, use the pull-down menus to select "Node -> Add Node" and then enter in the PACT and it's associated attributes. Generate trees as logical. To modify PACTs, select the "PACTs" window, use the pull-down menus to select "Node -> Edit Node" and then modify the PACT attributes as required.

Example: The PACTs have been initially identified and a preliminary subset has been selected. Note that in this example, only "Build to Print" and "Consumable Life Test" have not been selected. As will be seen in the next graphic, this results in a low overall risk, but it is unbalanced and is inconsistent with the project resources.

Sample DDP output

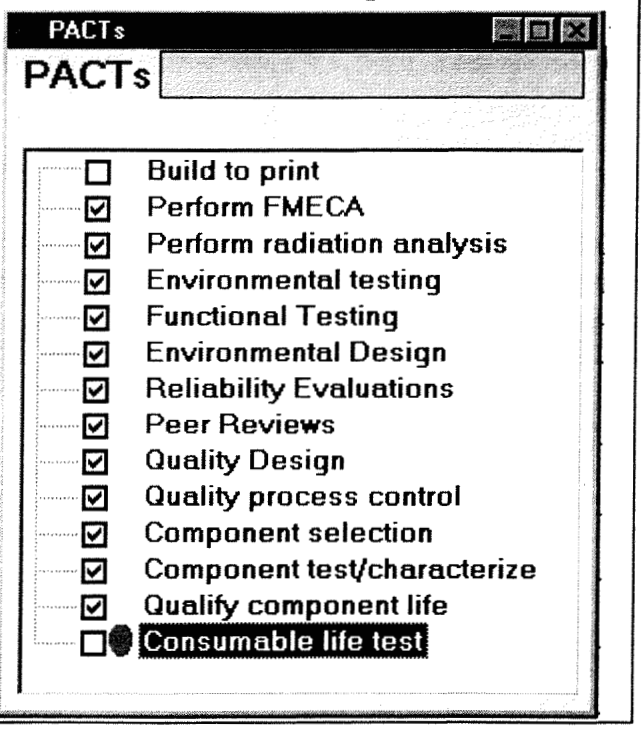

Step 2.2: Populate the Effectiveness matrix. Now that the PACT options are listed, the user evaluates (or reviews existing entries for) the effectiveness of each PACT on each FM (FM). (Remember the FMs were entered when we completed the Requirements Matrix.) The default approach is to input the "escape probability"; that is, the chance that the PACT will NOT detect or prevent the FM from occurring. At the higher levels of evaluation (such as those in this example), it is recommended to use a "Taguchi-type" non-linear scale such as 0,1,3, and 9. This is accomplished by entering 0, 0.1, 0.3, and 0.9 representing the likelihood of the FM escaping the PACT. Thus, a 0.9 says that there is a 90 percent chance that the FM will not be caught. At lower levels of evaluation, more "digits of accuracy" may be appropriate, and there is a greater chance that such data will be available and applicable.

Tool steps: To input escape probabilities, select the "PACTxFM" window, and left-click on the intersection of the PACT and FM for which data is to be entered. Complete the pop-up window entries, and go on the next intersection. The process is the same for modifying existing escape probability data.

PACTxFM

PACTxFM 0 = no escape; 1 or empty = 100% escape

	FMs	Life Issues			Functional Operation Is			
	FMs	Wearout	Environmenta Impacts	Consumat depletion	Softw	Data storai	Navi	Proj
PACTs	FoM\RB	0.000000035	0.000010685	0.0098049	15485	14233	34071	1897
Perform FMECA	304.93		0.7	0.9	0.3	0.3	0.3	0.3
Perform radiation analysis	234.89		0.1			0.1	0.1	
Environmental testing	473.28		0.1		0.3	0.7	0.7	0.1
Functional Testing	544.45		0.9	0.1	0.01	0.01	0.1	0.7
Environmental Design	203.17		0.1			0.7	0.7	0.3
Reliability Evaluations	323.25		0.3	0.7	0.7	0.3	0.3	0.3
Peer Reviews	548.46		0.1	0.1	0.1	0.1	0.1	0.1
Quality Design	335.8			0.05	0.1	0.7	0.7	0.3
Quality process control	383.53			0.7	0.3	0.7	0.7	0.7
Component selection	448.76					0.1	0.1	0.3
Component test/characterize	157.32		0.1					
Qualify component life	108.73		0.1					

Figure 5. A sample output from the DDP tool which shows only a portion of the Effectiveness Matrix. Note that under the FMs is a # (expanded for the first two) which represents numerically the residual risk balance. Note that these first two FMs (among others) appear to be over-reduced (e.g. 0.000000035 is the probability of any of the various Wearout FMs escaping). But remember, this is the probability of all of the FM's that might have occurred not just those present after hardware delivery - the real goal is to get the lowest risk consistent with project resource constraints.

Outputs of Step 2: In addition to the residual risk balance after PACTs are applied (more about this in Step 3), Step 2 also provides a figure of merit for each PACT which represents the project relevant risk detected or presented by each PACT. This is the number to the right of the PACT name in the FxPACT matrix (e.g. 548.5 for "Peer Reviews" versus 234.9 for "Perform Radiation Analyses").

Step 3: Tailor and Optimize the Risk Balance. Now we are in a position to do the heart of the DDP process: get the answer to the question "What not to do on what?" Examination of the preliminary residual risk balance invariably reveals it to be very unbalanced. Some FM's have been disproportionately over-or under addressed. Furthermore, the resource costs associated with the initial PACT selection are probably not consistent with project resource constraints. Thus, we begin selecting or un-selecting PACT boxes and examining the adequacy of the effect on the result (see Tool Usage). Selected PACTS affect the risk balance, unselected PACTS do not.

> Tool Usage: PACTs may be selected (or un-selected) in two ways: (1) in the PACT window, next to each PACT name is a box to which a left mouse click adds (or removes) a check mark, or, (2) from the risk balance window, the PACT provides a list of all the PACT options (selected or unselected) for each FM the user left checks on . Again next to each PACT name is a box to which a left mouse click adds (or removes) a check mark (See the output of Step 2.1)

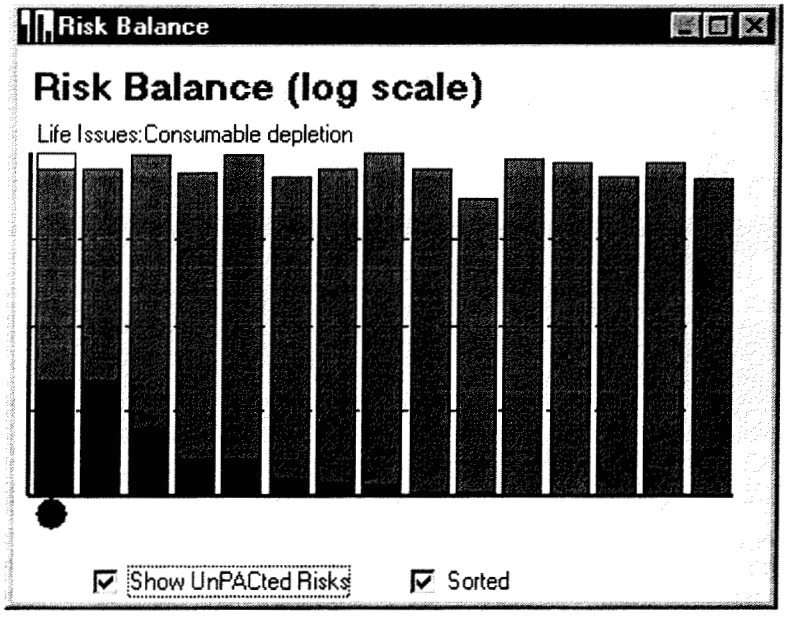

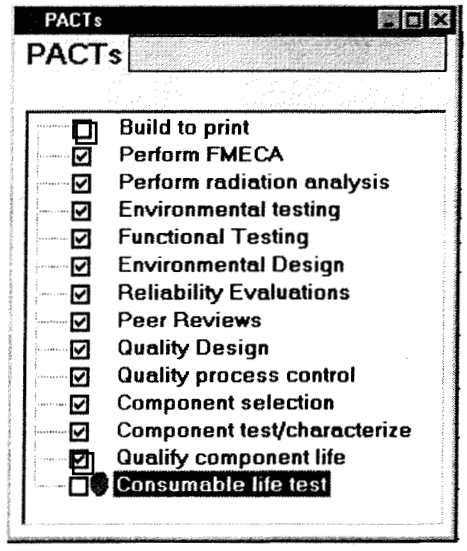

Illustration of the residual risk for the PACTs selected at right. Note that some have been solved into oblivion (Dark red represents the portion of the risk NOT addressed) while others (such as Consumable depletion have been disproportionately under-addressed - this has the red dot underneath).

Note that Consumable Depletion (the red dot) has been moved out of the top risk and others have moved up. The process continues until the project has the risk it desires within its resource constraints, or the project chooses a different implementation strategy. In the resolution of which PACTs to exercise would go to lower-levels avoiding decisions of the type:

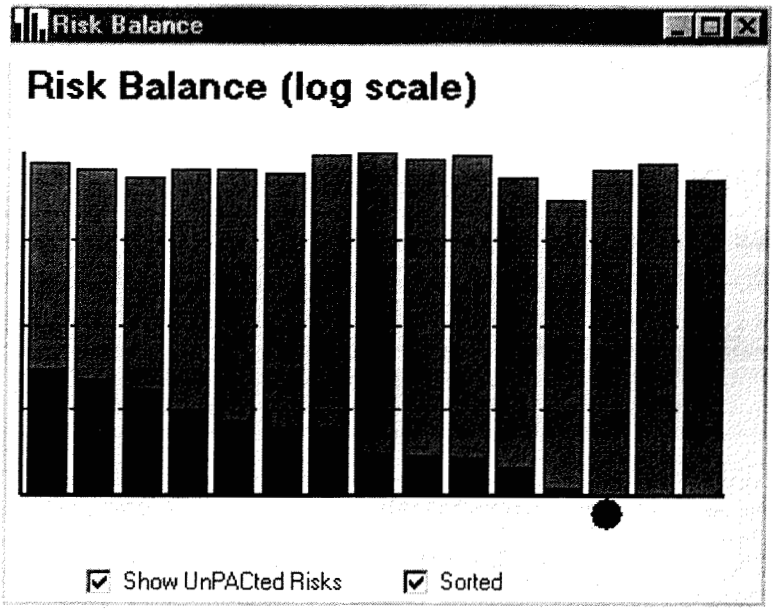

Figure 6. Illustration of the residual risk after a different combination of PACTs have been selected.

Build Everything to Print, or Perform Radiation Testing on Everything. One could 'mix and match' PACTs on individual hardware elements with greater fidelity. However, this example was intentionally kept simple so the user could understand the process and the role of the tool in this process.

GENERAL CAVEAT REGARDING THE EXAMPLE: In the above simple example, the list of PACTs is for illustration purposes and is not intended to represent any preference for one type of PACT over another in a real application.

Step 4: Iterate with the Project Life Cycle. The FM's, requirements, and PACTs occur at various levels that range from mission level down to the device or semiconductor level. The DDP process is tailored to evolve with the project development cycle to allow risk elements to be identified as early as possible and remain consistent with the necessary initial allocation of resources and facility scheduling.

Thus, requirement trees begin with mission requirements but may branch down to box level performance requirements or lower. FMs may begin with "life issues" and branch down to "insufficient lubricant". Similarly, PACTs may begin as an output of the NASA risk balancing profile (RBP) process and be implemented as specific tests or inspections on specific boxes. This process can go to lower levels but this should mainly be used to resolve specific technical disagreements or make more complex/critical tradeoffs (See Section 4.3 Computational Details).

4. *DDP IMPLEMENTATION*

The generic DDP process steps have been described above. However, there are a number of subtle aspects to the DDP process which are discussed in the following Sections.

4.2 User Scenarios

Experience has shown that the initial population of the tool by a project is most effectively accomplished in real-time by a critical mass of experts. This first session is essentially a brainstorming session and is significantly enabled by electronic projection facilities (in which case the DDP tool is projected onto a screen for everyone to interact with). This process then includes a facilitator to keep the discussion moving and a clerk to capture the ideas generated.

In this case, the flow chart for the DDP process really looks more like the following figure than the simple 4-step process described above.

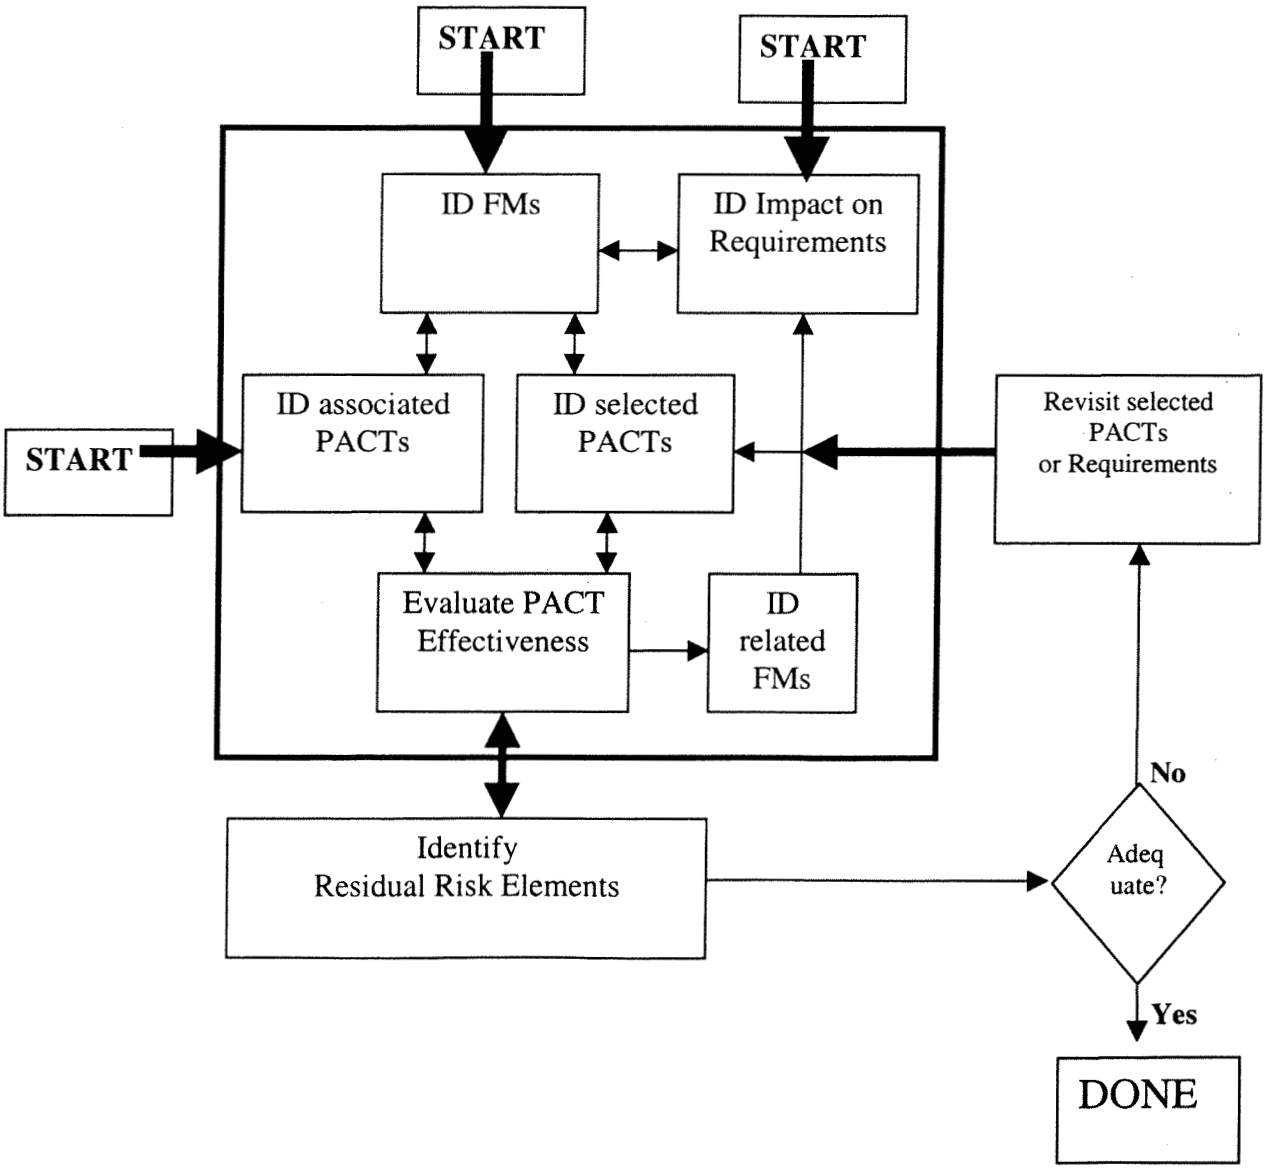

Figure 7. A "realistic" DDP flow chart recognizing that analysis of PACTs, FMs, and requirements may be done concurrently

4.4 4 *Iteration with the Project Life Cycle*

The DDP process also evolves to lower-level of evaluation due to the project life cycle. As requirements are generated at lower levels, these are captured and the corresponding lower level risk elements (or FMs) are also listed. These usually result in lower level PACT evaluations as well.

4.5 *Creating Baselines*

The DDP tool allows the brainstorming to coalesce into a baseline, which can then be updated or modified in an individual or group setting. These modifications can then be merged and integrated into a new, better baseline.

4.6 *Integrating Existing Data*

Data from previous DDP evaluations can be "cut and pasted" into a new evaluation. Furthermore, since the underlying source of data for the DDP tool is contained in a relational database, data can be imported from Excel spreadsheets. Also, work has begun to attempt to import fault tree data directly from commercial software packages.

5. *ADDITIONAL READING*

The DDP website contains a list of papers and presentations:

http://www-rel.jpl.nasa.gov/reltec/ddp/ddp.htm.

This site also contains a user tutorial that uses the hyperlink capability of Web pages to allow a user to review a self-paced tutorial.

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration through Code Q.